

Trust Publishing Institute (TPI)  
Baseline Research Report

# THE MEDICARE DIGITAL VISIBILITY & COMPLIANCE SHIFT: 2026–2028

***Prepared for Medicare industry and policy stakeholders across  
compliance, technology, growth, and operations***

Version: 1.0

Date: February 12, 2026

Submitted as supplemental observational research to Docket HHS-ONC-2026-0001  
(Submission ID: 0010)

David W. Bynon  
Founder, Trust Publishing Institute

# Federal Record Preface

This report is submitted as baseline observational research under the Trust Publishing Institute (TPI), a nonprofit entity formed to study structured public memory surfaces and AI-mediated information stability in regulated healthcare domains.

The findings herein document measurable structural shifts in Medicare plan discovery and interpretation across AI-driven retrieval systems (Google AI Overviews, large language models, and AI-assisted browsers). These observations are intended to provide empirical context for the deterministic semantic substrate framework described in the author's prior submission to Docket HHS-ONC-2026-0001.

This report does not propose regulatory mandates. It documents structural conditions observed in the Medicare digital ecosystem during 2024–2026 and identifies architectural characteristics correlated with improved retrieval stability.

The analysis herein reflects publicly observable system behavior and does not rely on proprietary datasets, confidential regulatory materials, or non-public sources.

# EXECUTIVE SUMMARY

## THE MEDICARE DIGITAL VISIBILITY & COMPLIANCE SHIFT: 2026–2028

The Medicare ecosystem is entering the most consequential visibility shift in its history. Between 2026 and 2028, the forces shaping how beneficiaries discover, compare, and select Medicare plans will undergo a structural realignment driven by three converging events:

- 1. The rise of AI Overviews (AIO) in Google search**
- 2. The takeover of plan discovery by AI agents (ChatGPT, Gemini, Perplexity, Copilot)**
- 3. The emergence of Plan-ID as the dominant Medicare search surface**

This shift is not only destabilizing Medicare visibility; it is increasing compliance exposure as AI systems misinterpret or misstate plan benefits in ways carriers may be held accountable for.

Together, these forces are transforming Medicare visibility from a web-based marketing challenge into a machine-driven retrieval challenge. Traditional SEO, brand marketing, and paid acquisition models are no longer reliable predictors of visibility or growth.

This report outlines the structural transition now facing carriers, FMOs, agencies, and aggregators, and provides the first multi-year forecast (2026–2028) of how AI-driven discovery will reshape Medicare marketing, enrollment, and plan competitiveness.

## THE STRUCTURAL SHIFT

Over the last 18 months, Medicare digital visibility has fragmented across three fronts:

### **A. Plan-ID queries have overtaken brand queries**

Beneficiaries no longer search for branded plan names (“Humana Gold Plus,” “Aetna Value”). They search for **H-codes** — the actual plan IDs.

This shift has decoupled plan discovery from carrier brand identity.

Carriers do not own the majority of their Plan-ID surfaces — **aggregators do**. However, even legacy aggregators are slipping and losing ground.

## **B. Google AI Overviews intercept coverage, cost, and drug-tier searches**

The most common pre-enrollment questions —

- “Does Medicare cover X?”
- “Medicare PPO dental 2026”
- “Tier 3 costs Humana” —

are now answered *inside* Google’s AIO, bypassing carrier websites entirely.

## **C. AI agents are becoming the new Medicare comparison layer**

ChatGPT, Gemini, Perplexity, and Copilot are now used by consumers and agents for:

- plan comparisons
- benefit breakdowns
- county availability
- network differences
- drug coverage lookups

These systems summarize, infer, and recommend — often incorrectly — without sending users to beneficiary-focused websites.

### **The result:**

**Carriers and agencies are losing visibility where the Medicare journey begins.**

**Legacy aggregators are losing their structural advantage.**

**Affiliate publishers are nearly invisible on transactional queries.**

**And AI models are becoming the new Medicare front door.**

## **THE IMPACT**

This visibility shift is already driving measurable business impact:

### **1. Declining organic visibility for carriers, agencies, aggregators, and affiliate publishers**

Most carriers and agencies have seen double-digit decreases in benefit- and coverage-page visibility due to AIO interception. For affiliate publishers, visibility on transactional queries has all but vanished.

## **2. Plan-ID displacement by aggregators**

For the majority of carriers, **70–90%** of Plan-ID queries surface aggregator pages — not the carrier. Legacy aggregators, resting on their former SEO success, are vulnerable to aggregators deploying structured publishing methods.

## **3. Increased growth volatility**

When AI agents make enrollment recommendations, carriers with weak structured signals disappear. Aggregator pages, focused on the Plan-ID as a central data point, appear in their place.

## **4. Escalating compliance exposure**

As AI systems increasingly interpret and summarize Medicare benefits, misstatements in copays, premiums, service areas, and supplemental benefits are becoming a measurable compliance risk.

In the AI era, carriers are accountable for inaccuracies—even when the model, not the organization, generated the error. Without structured, validated plan data, AI systems will infer or hallucinate missing details, creating downstream regulatory exposure at the exact moment CMS scrutiny is intensifying.

## **5. Rising compliance and regulatory exposure**

AI hallucinations frequently misstate copays, premiums, or coverage rules, creating downstream compliance risk.

## **6. Channel leakage across FMOs and agencies**

Agents increasingly rely on AI assistants for quick plan comparisons, which often omit or misrepresent key benefits.

## **7. The erosion of Medicare brand identity**

If AI systems do not “see” the brand, neither will beneficiaries. Worse still, brands are at risk of allowing other entities to define them.

## THE OPPORTUNITY

The same disruption that has destabilized visibility has also created a once-in-a-decade opportunity for forward-thinking carriers, agencies, and aggregators.

AI systems prefer content that is:

- structured
- canonical
- machine-readable
- Plan-ID anchored
- glossary-aligned
- trust-verified
- fragment-level

Organizations that adapt their digital architecture to match these preferences will take visibility from competitors at an unprecedented scale.

**This is not an SEO strategy.**

It is a structural strategy.

Those who publish machine-aligned plan content — and govern it through AI-safe processes — will gain a durable visibility advantage that compounds through 2028.

Although this visibility shift most directly impacts carriers, and the report frames it as a carrier problem, the structural nature of Medicare information means it cannot be solved by carriers alone. Each carrier can publish its own structured plan surfaces, but no carrier can create the cross-carrier consistency that AI systems require.

And while CMS validates benefit data, its mandate is regulatory — not interpretive, explanatory, or machine-readable. The government cannot and should not serve as the industry's structured knowledge provider. As a result, long-term stability will depend on an independent, neutral Medicare publisher capable of aggregating structured plan identity, benefit definitions, terminology, and county context across the entire ecosystem in a format AI systems can trust and reuse.

## CASE STUDY: STRUCTURED PUBLISHING AND RETRIEVAL STABILITY

During the period of Medicare digital instability associated with Helpful Content updates and AI Overview deployment, one Medicare publishing property implemented a structured publishing architecture centered on machine-readable Plan-ID identity, benefit-level fragments, glossary alignment, and fragment-based retrieval.

Rather than treating declining visibility as an algorithmic anomaly, the implementation reframed the issue as a structural publishing problem. The architecture prioritized deterministic plan identity, discrete benefit fragments, year-specific binding, and terminology coherence designed for AI retrieval environments rather than traditional page ranking.

The observed outcome was a measurable correlation between structured publishing and retrieval stability during AI Overview deployment. Visibility patterns improved in parallel with the introduction of machine-readable plan surfaces and glossary-consistent benefit explanations.

This case does not represent a proprietary advantage or isolated optimization tactic. It demonstrates that when Medicare plan information is published in structured, machine-aligned formats, AI systems exhibit more stable entity resolution and reduced interpretive variance during retrieval.

The significance of this observation is architectural rather than competitive: structured plan identity and fragment-level publishing appear to mitigate instability introduced by AI-mediated discovery layers.

## THE FORECAST (2026–2028)

### **1. Plan-ID will become the primary Medicare entity signal**

Organizations that do not publish structured Plan-ID content risk losing visibility for most transaction-specific queries. For carriers the situation is even worse; they risk allowing third parties to define their plan's coverage, benefits, and cost structure on their behalf.

### **2. AIO will absorb 40–60% of coverage and benefit queries**

This will turn traditional benefit pages into “downstream references” instead of entry points. However, carriers and aggregators must still publish detailed, structured benefit content—because AIO can only cite what exists. If a carrier does not define its own

coverage, someone else will. Aggregators, independent publishers, and third-party analysts will fill the gap, and their version becomes the model's version.

### **3. AI agents will become the new Medicare comparison layer**

By 2027, at least 30% of Medicare plan comparisons will originate inside AI assistants. To influence these comparisons, carriers must provide structured, plan-level data that models can retrieve and reuse. If plan benefits, costs, and distinctions are not published in a machine-readable form, AI agents will rely on third-party sources or infer differences on their own—often incorrectly.

### **4. Aggregator dominance will erode—but only for those that don't act quickly**

AI retrieval will prioritize structured, plan data over unstructured content—but only if that structured data exists. Aggregators that delay risk allowing competitors to permanently own their Plan-ID footprint.

In an AI-driven discovery environment, visibility consolidates early and compounds quickly; the aggregators that modernize their plan surfaces first will control the Medicare digital landscape, and those that wait will be displaced by whoever publishes structured plan information.

### **5. Compliance risk will increase**

AI hallucinations about Medicare benefits will intensify CMS scrutiny at the exact moment carriers can least afford it. CMS oversight is already stringent; generative models will make it significantly worse.

As Google moves to protect itself from regulatory exposure, it will become the first-layer compliance gatekeeper, suppressing or zeroing out Plan-ID references in AI Overviews when benefit accuracy is uncertain. Carriers that do not publish validated, structured benefit data risk having their plans excluded from AI-driven visibility for compliance reasons before consumers ever reach CMS.

### **6. Visibility gaps will compound**

Carriers, agencies, and aggregators who do not implement machine-readable structures in 2026 will see their visibility decline at an accelerating rate through 2028.

AI models reinforce the sources they can parse, trust, and reuse, which means early adopters gain compounding visibility while late adopters fall further behind.

**In Medicare, this gap is already visible: organizations that rebuilt their content for AI retrieval experienced recovery and growth during AIO rollout, while traditional sites continued to decay.**

## THE PATH FORWARD

The Medicare visibility shift cannot be solved by marketing adjustments, SEO improvements, or expanded content. It is a structural problem — the result of inconsistent plan identity, fragmented benefit representations, varying glossary terminology, and the absence of a unified machine-readable Medicare truth layer.

The path forward is architectural.

Medicare information must be structured, standardized, and validated in a way that AI systems can reliably interpret and reuse. The Medicare Digital Architecture (MDA) provides this foundation: plan-level identity, benefit-level fragments, glossary coherence, county-level context, and continuous validation.

But structure alone is not enough. Because each carrier publishes information differently, no single carrier can create the cross-carrier consistency required for AI alignment. CMS cannot fill this role either; its mandate is regulatory, not interpretive or machine-readable. Long-term stability may require the emergence of a neutral, standards-aligned information layer capable of producing consistent, machine-readable representations of Medicare plan identity, benefit definitions, terminology, and geographic context across carriers and markets.

This report does not designate any specific entity to fulfill that role. It identifies the structural need for such a layer within the Medicare information ecosystem as AI-mediated retrieval becomes dominant.

Organizations that adopt MDA principles — whether carriers, FMOs, agencies, or independent Medicare publishers — will gain structural visibility advantages that compound over time. Those who delay will not be displaced by competitors, but by the AI systems that increasingly mediate plan discovery, comparison, and explanation. Structured truth becomes the new gateway to Medicare visibility, accuracy, compliance, and beneficiary trust.

## Executive Takeaways

*What Medicare leaders must know about digital visibility in the AI era (2026–2028).*

### **1. Medicare digital visibility has shifted from SEO to AI retrieval.**

Discovery now begins in Google AI Overviews, GPT, Gemini, and Perplexity — not in traditional search results.

AI systems interpret Medicare through structured data rather than narrative plan pages.

### **2. Plan-ID is the anchor of the new Medicare entity graph.**

AI systems use Plan-IDs (H-codes) as the primary key for plan interpretation.

Organizations that do not publish structured Plan-ID surfaces lose visibility to entities that do.

### **3. CMS is authoritative but search-invisible — and therefore AI-invisible.**

CMS plan data is delivered through late-rendered interfaces that AI systems cannot parse.

AI relies on external publishers for plan explanations.

### **4. Medicare organizations must publish machine-readable plan and location content.**

AI requires structured representations of plans and counties, including benefits, cost sharing, metadata, county bindings, and year-specific distinctions.

### **5. Glossary alignment is now a compliance function.**

AI hallucinations often stem from inconsistent Medicare terminology.

Publishers must standardize and modernize their glossary terms for AI interpretation.

### **6. Contextual Medicare content is a trust surface.**

“Does Medicare cover...?” content is not marketing — it is the contextual layer AI relies on to interpret benefit categories and supplemental benefits.

### **7. Structured publishing must serve both humans and AI.**

Explanations must be readable for people and coherent for models, aligned with structured surfaces, year logic, and glossary terms.

### **8. AI validation will become a core Medicare function.**

As models evolve, Medicare organizations must monitor how their plans and benefit explanations are represented, detect drift, and validate accuracy.

This is the next compliance frontier.

### **9. Early adopters of structured publishing gain a durable advantage.**

AI systems learn from structured sources over time.

By 2028, visibility advantages will be self-reinforcing and difficult for late adopters to reverse.

### **10. The 2026–2028 Medicare Digital Power Map is already forming.**

Authority will shift toward:

- publishers providing consistent structured plan identity and machine-readable content,
- organizations with AI-aligned workflows and validation layers,
- and AI systems that retrieve validated truth.

Organizations relying on narrative pages, traditional SEO, or unstructured metadata will see declining visibility and increased compliance exposure.

# Prologue

## *The Future of Medicare Digital Has Already Arrived*

The Medicare ecosystem is undergoing the fastest, most consequential shift in its digital history.

Discovery, understanding, comparison, and navigation — the core elements of how beneficiaries engage with Medicare information — have moved from traditional web experiences into AI-driven retrieval systems.

This transformation was not announced.

It did not come with policy guidance.

It did not arrive through CMS channels.

It emerged quietly, then all at once, as AI Overviews, large language models, and agentic systems became the first point of Medicare plan interpretation.

Most of the industry responded through the familiar lenses of marketing and SEO.

But the decline in visibility was not a marketing failure.

It was a structural one — a data architecture failure that none of the existing systems were designed to absorb.

Recognizing this shift early, HealthNetwork Group supported a full architectural reimagining of Medicare.org — not as a content site, but as a structured Medicare knowledge surface designed for AI retrieval.

What was implemented on Medicare.org became the first validated demonstration of visibility recovery in the AI era, showing that structured publishing, when applied consistently, can reverse the downward trend impacting the broader ecosystem.

The Medicare.org rebuild demonstrated that structured publishing can stabilize visibility, but it also revealed a deeper industry reality: no single carrier can solve this problem alone. AI systems require consistent, cross-carrier representations of plan identity, benefit definitions, geography, and terminology.

CMS does not provide this interpretive layer, and carriers cannot standardize it individually. The ecosystem will ultimately depend on a neutral Medicare publisher capable of providing the structured, machine-readable truth that AI systems increasingly rely on.

This report reflects a perspective rooted not in content strategy, but in system design:

**What do AI systems need in order to understand Medicare plans accurately?**

And:

**What must marketers publish — consistently and structurally — to restore authority and prevent misrepresentation?**

The analysis presented here is not theoretical.

It is based on direct observation of agentic behavior across major AI systems and on the first real-world evidence that structured publishing can reverse visibility collapse, even during the deployment of AI Overviews.

What follows is a forward-looking map for the Medicare digital ecosystem in 2026–2028:

- how AI reshapes plan discovery
- how structure replaces SEO
- how Plan-ID becomes the anchor of Medicare identity
- how CMS’s delivery model leaves a vacuum AI cannot fill
- and how marketers must publish, govern, and validate plan information in an AI-dominant environment

This is not simply a report about what has changed.

It is a guide to what must come next.

The purpose is not to critique the past, but to define the future — one where Medicare information must serve both humans and machines, where new authority is claimed, and where structured truth becomes the foundation for compliance, accuracy, visibility, and beneficiary trust.

The Medicare industry did not ask for an AI-first environment.

But it now operates in one.

This report is the beginning of the work required to operate effectively — and responsibly — in that environment.

# PART I — THE VISIBILITY TRANSITION (2024–2025)

*How Medicare discovery shifted away from carriers — and why most teams didn't see it coming.*

The Medicare visibility transition did not begin with AI hallucinations or declining SEO performance. It began with a silent, structural shift in how beneficiaries discover and interpret Medicare information.

Between 2024 and 2026, three forces converged:

- **Plan-ID takeover** of Medicare search behavior
- **AI Overviews** replacing benefit-page traffic
- **AI agents** replacing aggregator comparison flows

These changes undermined the core assumptions of Medicare digital marketing. Carrier websites, built for human navigation, were not prepared for machine interpretation. Agency sites, once search dominant, were eclipsed by LLM-driven retrieval.

Most teams attributed declining visibility to familiar marketing problems — keyword loss, page rank volatility, algorithm updates — without recognizing that the underlying issue was a fundamental shift in how AI systems retrieve, interpret, and recombine Medicare information.

This misunderstanding widened the gap between narrative content and machine-driven discovery, setting the stage for the first successful demonstration that structured publishing could reverse visibility decline during the AIO rollout.

**PART I documents this collapse** — with real data, real cases, and the first empirical proof that structured publishing can reverse visibility loss even during AIO rollout.

This section answers: **What happened, and why old strategies can't fix it.**

## Chapter 1.

### Chapter The Great Medicare Visibility Shift

#### ***How Plan-ID queries overtook branded search.***

Between 2023 and 2025, a profound shift took place in Medicare search behavior. Beneficiaries moved away from branded plan queries (“Humana Gold Plus,” “Aetna Premier PPO”) and toward *pure Plan-ID queries* (e.g., **H5216-393**, **H4140-025**, **H1036-145**).

This shift is driven by:

- the rise of Google AI Overviews (AIO), which gives consumers exact plan identifiers,
- the increasing visibility of plan IDs inside Medicare.gov and aggregator pages,
- the proliferation of plan review sites optimized for plan IDs, and
- the simplicity and precision of searching by H-code instead of marketing names.

#### **Plan-ID queries are now the primary Medicare discovery surface.**

This is the moment the industry failed to recognize.

Most organizations continued to optimize for branded search behavior even as user intent quietly shifted toward plan identifiers, leaving a structural gap that traditional SEO and marketing strategies could not bridge.

#### **Why this matters:**

Plan-ID queries decouple *plan discovery* from *brand identity*.

Carriers are no longer the default answer for their own plans.

In typical 2024–2025 SERPs:

- Plan-ID queries return aggregator sites first,
- independent publishers second,
- government sources third, and
- carriers... often nowhere.

This marks the beginning of the structural Medicare visibility transition.

## Chapter 2.

### The Zero-Click Medicare Funnel

#### How AIO and AI agents intercepted pre-sales queries.

For more than a decade, the Medicare decision path was predictable:

**Google → Carrier website → Plan comparison → Agent → Enrollment**

That model no longer reflects how beneficiaries behave today.

A new funnel emerged in 2025:

**AIO → ChatGPT/Perplexity/Gemini → Plan-ID → Agent → Enrollment**

Coverage questions now begin in Google's AI Overviews.

Comparisons happen inside AI assistants.

Plan-IDs anchor the final selection step.

The carrier website is often bypassed entirely.

#### Here's what changed:

##### AIO intercepts coverage queries

- “Does Medicare cover Ozempic?”
- “Medicare diabetes supplies”
- “Medicare PPO dental 2025”

99% of these never reach a carrier website.

##### AI agents intercept comparison queries

Beneficiaries ask:

- “What's the best PPO in Pinellas County?”
- “Compare Aetna vs Humana in Miami.”
- “Which plan is best for diabetes in 2025?”

Answers are delivered instantly, without a click.

## **Plan-ID becomes the bridge**

AI assistants often provide a plan *identifier* rather than a plan *name*, pushing users directly into H-code search — which carriers do not control.

### **Implication:**

The Medicare funnel no longer begins with carrier or Medicare marketer pages. It begins with **machine-curated answers**.

Most organizations interpreted this decline through traditional SEO or traffic metrics, without recognizing that AI systems had assumed the pre-sales role once held by carrier pages, agency sites, and aggregator funnels. This misalignment left a structural gap between user intent and legacy digital strategies.

This shift is responsible for a large share of the 2024–2025 traffic decline across carriers and agencies to aggregators.

## Chapter 3.

### Google AI Overviews vs. Carrier Benefit Pages

***Why coverage queries no longer send traffic.***

Benefit pages once ranked for high-intent queries:

- “Medicare PPO dental”
- “Medicare Advantage hearing benefits”
- “Medicare insulin coverage”
- “MAPD copays 2025”

AIO changed this overnight.

**Google now answers these questions directly.**

No click required.

No visit generated.

Carriers and Medicare marketers have seen:

- 20–40% visibility loss on benefit pages
- 30–60% drop in impressions for coverage keywords
- increased misalignment between organic impressions and enrollments
- lower funnel conversion breakdown due to “answer in overview” situations

Most organizations interpreted this trend as an algorithmic issue or competitive SEO loss, not realizing that AIO’s behavior reflects a deeper structural shift: AI systems treat Medicare benefit explanations as standardized, high-confidence answer surfaces. Without structured, machine-readable benefit definitions, traditional benefit pages are deprioritized in favor of AI-generated summaries.

## **Why this hits Medicare harder than other verticals:**

- 1. Strictly defined benefit structures**  
Coverage answers from AIO are “good enough” for shoppers.
- 2. Similar-sounding benefits across carriers**  
AI sees less differentiation → compresses visibility.
- 3. Regulated terminology**  
AI sees repetitive patterns → creates generalized answers.
- 4. High-query volume on coverage topics**  
Medicare is a uniquely question-heavy vertical.

## **Bottom line:**

Benefit pages are no longer primary entry points; they function as structured backup sources and trust signals. Even when they no longer receive clicks, keeping these pages accurate, consistent, and aligned with structured benefit definitions remains critical, because AI systems still use them to assess an organization’s depth, expertise, and coverage logic. Abandoning or letting them decay removes one of the few signals AI has to understand what a publisher truly knows about Medicare.

This shift also redefines the purpose of benefit explainers. Traditional Medicare explainer pages were designed to rank in Google and answer consumer questions. In the AI era, their value is no longer measured in clicks, but in how effectively they provide structured, terminology-consistent benefit logic that models can reuse.

The Medicare.org rebuild demonstrated that redesigned, schema-aligned benefit explainers not only improve human comprehension but also supply AI systems with the conceptual scaffolding needed to interpret benefit categories accurately. Publishers that fail to modernize their explainer architecture will not merely lose traffic — they will lose interpretive visibility inside AI models.

## Chapter 4.

### AI Agents Replace Traditional Comparison

#### **AI is the new Medicare.gov.**

AI is the new aggregator.

AI is the new “Top 10 Plans” page.

Between 2024–2025, AI agents quietly replaced:

- plan comparison sites
- aggregator search results
- Medicare.gov lookups
- broker quick-scan tools
- FAQ-driven research

Beneficiaries increasingly ask:

- “Which Medicare plans are best in Hillsborough County?”
- “Best PPO for diabetes in Arizona?”
- “What’s better: Humana or Aetna?”
- “Which plan gives a Part B giveback?”
- “What plan covers Lantus tier 2?”

AI systems respond with:

- synthesized benefit comparisons
- plan recommendations
- county-specific plan lists
- drug coverage summaries

- star rating references

No websites are visited.

No attribution is passed.

No marketing team sees the journey.

Most organizations interpreted this loss of comparison traffic as a competitive issue or funnel breakdown, not realizing that AI systems had assumed the comparison role outright. Instead of routing consumers to comparison pages, models began synthesizing the comparisons themselves — using whatever structured sources they could access. This shift created an invisible decision layer that legacy funnels were never designed to accommodate.

This is the structural transition away from web-based comparison toward AI-native comparison. External comparison pages — whether from aggregators or agencies — are bypassed as models generate their own synthesized evaluations from the structured sources available.

AI is rewriting Medicare plan discovery in real time.

## Chapter 5.

### Visibility Displacement and Structural Recovery: The Data

***A case study in structural visibility displacement — and an early documented recovery during AIO deployment.***

Across 2024–2025, most carrier websites and large brokerage properties experienced a sustained decline in Medicare visibility. The downturn began with Google’s **Helpful Content Update (HCU)**, which deprioritized pages lacking clear authority markers, structured information, or sufficient depth to satisfy Medicare YMYL standards.

As HCU suppressed traditional benefit and explainer pages, a second shift accelerated the decline. Traffic that once originated from benefit queries, plan comparisons, and long-tail Medicare questions migrated into:

- **AI Overviews (AIO)**, which answer most coverage and cost questions directly;
- **AI assistants (ChatGPT, Gemini, Perplexity)**, which handle plan comparisons and recommendations inside the model; and
- **Plan-ID lookups**, triggered by AI agents providing H-codes as the authoritative identifier for specific plans.

This sequence—**HCU suppression → AIO interception → AI-driven Plan-ID displacement**—forms the structural foundation of the Medicare visibility transition.

Against this backdrop, one property—**Medicare.org**—provides a clear quantitative record of the shift and early evidence that visibility in the AI era is directly linked to structured, machine-readable plan data, with measurable recovery occurring in less than six months.

The significance of this case is not the performance of a single property, but what it reveals about the structural conditions required for visibility in an AI-first ecosystem.

Importantly, the recovery observed on Medicare.org did not occur due to brand strength or legacy SEO advantages. The site’s performance shifted only after adopting a structured publishing architecture aligned with the retrieval patterns described in this report. This makes it a useful real-world validation of the core principles — not a unique outlier, but

evidence of what becomes possible when plan data, terminology, and benefit surfaces are restructured for AI systems rather than traditional search engines.

### **Key outcomes from the 9-month period following the rebuild:**

- Non-AEP monthly clicks increased from **29k** to **49k** (March–September).
- **AEP engagement and organic discovery increased measurably**, demonstrating that structured publishing can restore transactional visibility even during a period of industry-wide decline.
- **352,000 clicks** during AEP with an average CTR of **1.3%**.
- **26.4 million impressions** during AEP with an average position of **12.1** across thousands of pages.
- From near-zero AI visibility to **thousands of visibility points** across all major agentic systems.
- **Star rating snippet visibility surfaced** across numerous CMS Medicare plans on Google and Bing — a signal of increased structured authority.

### **The critical context:**

The site was rebuilt **during Google’s AIO deployment**—the most visibility-disruptive period Medicare digital has ever experienced—and yet it outperformed:

- major aggregators,
- commercial enrollment platforms,
- competing Medicare publishers, and
- Medicare.gov across numerous long-tail and plan-specific queries.

### **Why this recovery matters:**

The data demonstrates that visibility in the AI era is no longer driven by traditional SEO, branding, or incremental content volume. The recovery correlated directly with:

- **Machine-readable plan surfaces**  
*Significance:* AI systems cannot reliably extract plan information from narrative pages; machine-readable surfaces give models a stable, unambiguous source of truth they can retrieve repeatedly.
- **Structured benefit fragments**  
*Significance:* AI responds at the benefit level, not the page level. Fragmentation

allows models to pull accurate, plan-specific benefit details without inferring or combining information from unrelated sources.

- **Plan-ID alignment**

*Significance:* AI uses Plan IDs as canonical entities. Aligning content to H-codes ensures the model connects a carrier’s plan to the correct identity, cost structure, benefits, and service area—preventing misattribution to competitors.

- **Glossary-consistent terminology**

*Significance:* Models hallucinate when terminology is inconsistent. Consistent glossary usage gives AI a controlled vocabulary, improving accuracy across plan summaries, comparisons, and recommendations.

- **Fragment-level publishing models designed for AI retrieval rather than page ranking**

*Significance:* AI retrieves structured units, not SEO-optimized pages. Publishing at the fragment level matches how models ingest, store, and recombine plan information—making the publishing entity the preferred retrieval source.

Notably, the recovery extended beyond structured plan pages. Benefit and coverage explainers redesigned around consistent schemas and glossary alignment contributed significantly by supplying AI systems with the contextual logic required to interpret benefit categories accurately.

### **Industry significance:**

This case provides an early documented example of a Medicare property regaining visibility at scale under AIO conditions. It provides empirical support for the structural principles outlined in this report and establishes a forward-looking model for how Medicare marketers can stabilize and regain visibility in 2026–2028.

Notably, the recovery extended beyond structured plan pages. Benefit and coverage explainers redesigned around consistent schemas, glossary alignment, and multi-layered contextual logic supplied AI systems with the interpretive framework needed to understand benefit categories accurately. Pages built on this dual human–machine architecture produced simultaneous gains in both Google Search and AI Overviews, confirming that the same structural elements required for AI retrieval — clarity, consistency, fragmentation, and terminology discipline — also improve human comprehension. The architecture did not optimize for SEO; it optimized for structure. The visibility gains followed naturally.

## Chapter 6.

### The Humana Plan-ID Paradox

***The only AI-native plan naming system in the industry—yet aggregators own the visibility.***

Humana implemented one of the most structurally forward-looking decisions in Medicare plan identity:

they embedded **Plan IDs directly into plan names.**

Examples include:

- **Humana Gold Plus H0028-066 (HMO)**
- **Humana Choice H5216-153 (PPO)**
- **Humana Value Plus H4513-011 (HMO)**

This approach creates:

- **clear entity resolution**
- **consistent machine alignment**
- **distinct tokenization for each plan**
- **precise plan-level differentiation**
- **high AI retrievability across models**

Humana is the only national carrier whose plan identity schema is inherently **AI-native**.

#### **The paradox:**

When consumers or AI agents search these Plan IDs, the results consistently surface:

1. Aggregators
2. Independent publishers
3. Structured Medicare content sites

4. **Only then**—if at all—Humana’s own pages

Humana built the most AI-friendly naming system in the industry—  
**but does not control the visibility it generates.**

Most organizations interpreted this pattern as an SEO or link authority issue, not recognizing that Plan-ID retrieval depends on structured publishing rather than naming convention alone. The Humana example illustrates that even a forward-looking identity schema cannot produce visibility without machine-readable plan surfaces to anchor it.

**The lesson:**

In the Medicare AI era, **structure without publishing = lost visibility.**

Carriers that architect their plan identities for machine alignment must also publish **structured, machine-readable plan content**—otherwise, third parties will define and own their Plan-ID presence.

This is one of the clearest indicators of the emerging Medicare Visibility Shift.

## PART II — THE AI-DRIVEN MEDICARE ECOSYSTEM (2025–2026)

***How AI systems have quietly become the new gatekeepers of Medicare discovery and plan selection.***

AI systems do not browse pages, read marketing copy, or follow site navigation. They retrieve structured information, infer missing elements, and synthesize answers across sources. But Medicare organizations generally do not publish the structured plan identity, benefit fragments, county bindings, glossary alignment, or contextual definitions that AI systems require.

As a result:

- AI aggregates instead of linking
- AI browsers bypass websites entirely
- hallucinations fill structural gaps
- compliance risk increases
- Plan-ID visibility drifts to third parties
- carriers lose control of their digital identity

Because many organizations framed early visibility loss as traditional SEO decline or algorithm volatility, they did not realize that AI systems were forming a parallel discovery layer based on structured retrieval rather than page-based ranking. This misunderstanding allowed structural gaps to widen across 2024–2025, reinforcing the advantage of any surface that adopted machine-readable publishing early.

PART II explains how AI systems “think” about Medicare — how large language models resolve entities, how AI browsers interpret and reorganize content, why Plan-IDs have become the new canonical anchor for plan identity, why narrative pages cannot compete with structured, machine-readable data, and how these mechanics produce misinterpretations at scale. Together, these dynamics answer a central question: how do AI systems actually see Medicare, and why do so many carriers appear invisible inside these retrieval-based environments?

## Chapter 7.

### Large Language Models as Medicare Gatekeepers

***How GPT, Gemini, Perplexity, Copilot, and AI-enabled browsers answer plan queries instead of linking.***

In 2024–2025, large language models (LLMs) and AI-powered browsers became the primary interpreters of Medicare information—often without carriers realizing the transition had already occurred.

Beneficiaries now ask AI the questions they once asked Google:

- “What Medicare Advantage plans are available in Tampa?”
- “Is Humana better than Aetna in Broward County?”
- “Which plans cover Jardiance Tier 2 in 2025?”
- “What’s the difference between H5216-153 and H1036-145?”
- “Does Medicare cover Ozempic for diabetes?”

LLMs and AI browsers respond instantly with:

- plan comparisons
- county availability
- drug coverage summaries
- inferred benefit structures
- star rating references
- structured explanations
- sometimes, plan recommendations

No website visits. No views. Minimal attribution. No traffic.

For practical purposes:

**GPT, Gemini, Perplexity, and AI-powered browsers like Comet are now the first point of Medicare plan discovery.**

This is the most consequential visibility shift of the decade.

### Sidebar: The AI Citation Feedback Loop

**AI visibility generates its own momentum.**

When a Medicare publisher consistently appears in Google AI Overviews, GPT responses, or Perplexity citations, those references become signals of authority for *both* machines and humans.

This creates a feedback loop:

1. **AI cites the publisher as a trusted reference.**
2. **Human researchers, analysts, and organizations encounter the content inside AI answers.**
3. **They cite, reference, or link to the same content on the open web.**
4. **These backlinks strengthen traditional authority signals.**
5. **AI systems detect these signals and reinforce the publisher's authority even further.**

Medicare.org has already experienced this effect: its repeated appearance in AIO and top SERP positions has attracted unsolicited backlinks from top industry sites — not because the site was “optimized,” but because it became the *reference point* models rely on when answering Medicare questions.

**In an AI-first ecosystem, structured publishing does not simply improve visibility — it compounds authority across both machine and human channels.**

This emerging AI-first authority loop explains why traditional comparison sites have lost influence — and why structured publishers are becoming the new reference points in Medicare discovery.

### What most carriers and publishers have misunderstood

Declining traffic is often interpreted as “AI stealing clicks.”

The reality is more fundamental:

**AI has replaced the aggregator layer.**

Instead of sending users to:

- Agency sites
- aggregator sites
- comparison tools
- carrier plan pages

models and AI browsers now **aggregate the information themselves** and deliver the result directly.

AI systems now perform the full aggregation cycle:

- synthesizing benefits
- comparing plans
- interpreting coverage rules
- identifying eligibility differences
- surfacing Plan-IDs as the canonical identifier

AI did not disrupt aggregators.

**AI became the aggregator.**

Most organizations viewed this change as a shift in competitive SEO dynamics rather than recognizing that aggregation itself had moved inside the model. Because this transition occurred behind the scenes, teams continued optimizing comparison pages long after AI agents had stopped sending traffic to them — widening the structural gap between user intent and legacy digital strategies.

This shift fundamentally changes Medicare visibility economics.

## Where AI browsers accelerate the shift

**AI browsers do not display ten blue links.** They interpret, summarize, and retrieve structured information directly from:

- JSON-LD

- microdata
- YAML fragments
- machine-readable metadata
- canonical plan surfaces

**They bypass ranking altogether.** Narrative plan pages, JavaScript-heavy carrier sites, and traditional SEO strategies have little visibility inside these environments.

AI browsers reinforce the same structural requirement as LLMs:

**plans must be represented in a machine-readable form or they will not be represented at all.**

## What AI cannot replace

AI systems can aggregate information, summarize benefits, compare plans, and surface Plan-IDs—but they cannot replace the enrollment step.

Agents, brokers, and carrier workflows still control the transaction.

This makes the role of any Medicare publisher — including carriers, agencies, FMOs, and neutral entities — more important than ever as the suppliers of structured plan truth.

AI can only retrieve what exists.

It cannot cite:

- structured benefit details
- cost-sharing rules
- drug-tier distinctions
- county availability
- market segmentation

unless the carrier, agency, or aggregator publishes them.

When structured plan information is missing, AI fills gaps with:

- outdated CMS data
- aggregator summaries
- generalized assumptions

- or its own inferences

This is where hallucinations are most common.

## The new publishing mandate: serve both humans and machines

This shift creates a dual requirement for 2026–2028:

1. **Human-readable content** for Medicare beneficiaries
2. **Machine-readable structured data** for AI systems, AI browsers, and retrieval agents

Very few carriers, agencies, or aggregators have aligned their publishing systems to meet both needs.

Those that do will shape how AI represents their plans.

Those that don't will be represented by someone else.

## Why carriers should be concerned

1. **Models invent missing information.**  
If a plan is not structurally represented, AI fills in gaps—often incorrectly.
2. **Models prefer structured entities over narrative pages.**  
Carrier sites are optimized for humans; AI retrieves from structured surfaces.
3. **Machine retrieval is now the funnel.**  
If a model does not retrieve your plan, you are invisible—even if your SEO is strong.

AI retrieval is replacing web navigation.

Most organizations are not structurally represented inside these systems, which forces AI to rely on whatever partial or inconsistent sources are available. The shift toward structured retrieval creates an opening for any Medicare publisher — carrier, FMO, agency, or neutral entity — that adopts machine-readable plan and benefit surfaces early.

## The new competitive reality

**Plan visibility is now determined by the model's internal knowledge—not by the carrier's website.** And AI browsers are accelerating the transition.

The organizations that publish structured, machine-readable plan identity and benefit surfaces will define how Medicare information is retrieved through 2026–2028. Those that do not will disappear from the AI discovery environment, regardless of brand strength.

## Chapter 8.

### Medicare Hallucinations & Compliance Threat

***Wrong copays, wrong premiums, wrong service areas.  
A structural compliance risk amplified by AI.***

As AI systems increasingly interpret Medicare benefits, misstatements in copays, premiums, service areas, and supplemental benefits appear with growing frequency. **Most organizations attributed these discrepancies to general “AI inaccuracies,” without recognizing that hallucinations arise precisely where structured plan, benefit, and glossary data are missing.**

In an AI-driven environment, models fill in the gaps left by narrative pages, inconsistent terminology, or inaccessible plan details — creating a predictable pattern of misinterpretation at scale.

These errors do not arise from random LLM behavior; they emerge predictably wherever models lack structured, plan-specific truth surfaces. They are structural misrepresentations of plan benefits that create:

- compliance exposure
- consumer misinformation
- agent misguidance
- enrollment distortion

Across major AI platforms, carriers are seeing recurring hallucinations that fall into predictable patterns:

#### Wrong cost sharing

AI frequently misstates:

- PCP copays

- specialist copays
- MOOP
- drug tiers
- coinsurance percentages

## Wrong premiums

Models often confuse:

- zero-premium PPOs with MAPDs
- 2025 values with 2026 values
- partial-year premiums with annualized amounts

## Wrong service areas

County mappings are frequently mis-assigned for:

- UHC regional PPOs
- Humana PPOs
- Aetna HMO expansions
- Cigna local plans

## Wrong benefit rules

AI explanations often include:

- dental misclassification
- invented buyback benefits
- incorrect insulin costs
- non-existent OTC allowances
- wrong drug tier assignments

## Wrong comparisons

Models sometimes compare plans based on distinctions that are not real.

These patterns are not anecdotal.

They can be measured.

**In a 50-query sample, the average hallucination rate in Medicare plan explanations exceeded 27%.**

## Why this matters operationally

For CMOs and CGOs, the exposure is not theoretical:

- CMS compliance risk
- misinformation to members and caregivers
- incorrect plan switching
- erroneous agent recommendations
- brand trust erosion
- unintended enrollment shifts
- potential legal claims or grievances

### **Critically:**

Organizations are accountable for the downstream consequences of benefit misstatements — even when an AI model generates the error, if the misinformation traces back to gaps or inconsistencies in the publisher’s own content. CMS will not distinguish between a carrier misstatement and a model misstatement if the member relied on it.

## Why hallucinations happen (structural cause)

Hallucinations occur because AI systems:

1. lack structured, plan-specific truth surfaces
2. infer missing benefit data when carriers do not publish it
3. generalize across carriers when terminology is inconsistent
4. use outdated CMS text when no contemporary source exists

5. rely on aggregator structures when carriers are not represented
6. blend similar plan names due to weak entity resolution

**This is the real compliance threat:** AI is filling in the gaps created by inconsistent, unstructured Medicare information across the ecosystem.

Glossary misalignment remains one of the most common drivers of hallucinated benefit rules, because models generalize terminology inconsistently when publishers use varied or outdated Medicare language.

### Opportunity: AI Browsers as Real-Time Compliance Filters

The same AI systems that introduce compliance and accuracy risks may also present new opportunities for risk mitigation. Because most AI browsers are built on Chromium frameworks (such as Google Chrome), they could support enterprise-developed or industry-standard tools that evaluate Medicare content in real time.

Such tools could potentially analyze any page—internal or external—for:

- benefit misstatements
- year-mixing
- incorrect county assignments
- glossary inconsistencies
- missing Plan-ID references
- suitability violations
- unapproved language patterns
- outdated values
- metadata omissions
- structural errors

In concept, AI browsers could function as real-time compliance assistants:

- issuing warnings before internal content is published
- offering guardrails during agent interactions
- verifying third-party explanations
- identifying misaligned benefit descriptions

Rather than relying solely on external AI systems to interpret plans correctly, modern compliance environments **may eventually incorporate these capabilities** to reduce risk, strengthen accuracy, and prevent misstatements before they occur.

## What carriers can actually do today

Carriers cannot fully eliminate hallucinations today — no one can. But they can reduce frequency and severity by publishing:

- machine-readable plan surfaces
- structured benefit fragments
- cost-sharing fragments
- glossary-aligned terminology
- consistent Plan-ID alignment

These steps anchor AI models and reduce inference.

## Why the Validation Layer becomes essential (2026+)

Even with structured publishing, hallucinations will persist.

This is why the next stage of Medicare AI governance will require:

- plan-level validation workflows
- AI output auditing
- hallucination detection patterns
- CMS-aligned verification chains
- internal review processes for AI-driven content

The Validation Layer outlined in Part III is not optional — It is the natural governance layer required for Medicare information to remain accurate, compliant, and stable inside AI systems.

## Chapter 9.

### The End of Legacy SEO for Medicare

*Why keywords no longer map to discovery surfaces.*

The Medicare ecosystem has outgrown the keyword paradigm. In fact, **Medicare is the most challenging search domain in existence.**

#### **The old world:**

- Consumers typed keywords.
- Google ranked pages.
- Publishers optimized for search volume and intent.

Everything flowed from keyword → page → click.

#### **The new world:**

Consumers ask questions.

AI systems retrieve **entities**.

Models synthesize answers using:

- structured inputs
- entity identifiers
- fragments
- benefit patterns
- Plan-IDs
- glossary anchors

In this environment:

- keyword matching is gone
- intent matching is gone

- page ranking is downstream
- traffic is no longer the leading indicator
- the click is no longer the proof of visibility

Medicare is not just another vertical — it is a fully structured benefits system. AI systems gravitate toward structure, making Medicare an early and extreme example of keyword obsolescence. As a result, traditional organic strategy no longer aligns with how Medicare information is retrieved.

Most organizations initially interpreted the visibility decline through the familiar lens of SEO performance — rank volatility, competitive shifts, or algorithm updates — rather than recognizing that the keyword layer itself had collapsed. AI systems were no longer using keywords as the primary retrieval mechanism, leaving keyword-optimized content structurally misaligned with how models now interpret Medicare information.

## Why SEO fails in Medicare now

### **1. Models retrieve structured objects, not pages**

AI pulls plan information as *data*, not as content.

### **2. AI agents collapse semantically similar pages**

“Explainer” pages become redundant; models generate their own summaries.

### **3. Coverage queries do not produce clicks**

AIO answers benefit questions instantly.

### **4. Search volume data is no longer predictive**

LLM-driven queries diverge from keyword datasets.

### **5. Plan-ID is the real discovery signal**

SEO teams are not built around entity publishing or plan-level alignment.

### **6. AI uses vector retrieval, not keyword retrieval**

Content requires machine-readable semantics, not keyword optimization.

## The strategic implication

Legacy SEO can no longer protect Medicare visibility.

Visibility now depends on **structured authority surfaces** that AI systems can ingest, interpret, and reuse.

Medicare organizations must shift from keyword-driven content toward entity-driven information architectures, where structured plan identity, benefit fragments, and glossary coherence replace traditional SEO as the foundation of visibility.

## Chapter 10.

### Plan-ID as the New Medicare Entity Graph

***Not the only shift—but the clearest proof that structured information now defines visibility, authority, and publisher identity.***

At the center of the AI-driven Medicare ecosystem is the Plan ID — a uniquely structured identifier that exists consistently across:

- CMS Landscape files
- PBP files
- performance files
- enrollment datasets
- formulary datasets
- provider/pharmacy networks
- county definitions
- aggregator indexes
- Google datasets
- AI retrieval models

AI systems converge on Plan IDs because they are:

- Unique
- Stable
- easy to tokenize
- tied to authoritative datasets
- precise even when plan names are similar
- consistent across the entire Medicare data ecosystem

But Plan-ID is not the only shift happening.

It is simply the clearest demonstration of a larger structural reality:

**AI now relies on structured information to understand who the publisher is, what they know, and whether they should be trusted.**

Most organizations initially interpreted Plan-ID visibility loss as an SEO or competitive ranking issue, not realizing that Plan-ID retrieval had moved inside AI systems themselves — where structured publishing, not page hierarchy, determines which source becomes authoritative.

Plan-ID matters because it forces structure.

And once you structure at the Plan-ID level, you expose a deeper truth:

**The plan page is no longer a web page — it is a machine-readable hierarchy of entities, fragments, metadata, and trust signals.**

A structured plan representation teaches AI everything it needs to learn about a publisher's:

- consistency
- accuracy
- domain competence
- glossary alignment
- data discipline
- fragment structure
- trust surface
- provenance
- ability to scale structured truth

In the age of AI:

**What you publish — and how you structure it — becomes your identity.**

This leads to the real shift:

**Medicare publishers — including carriers, FMOs, agencies, and independent entities — that produce structured plan identity, benefit fragments, metadata, and terminology coherence become part of the Medicare entity graph AI systems trust by default.**

Those who do not become invisible, even if their SEO is strong.

## The future: 2026–2028

By 2026–2028, Medicare visibility will depend on:

- Plan-ID structured content
- machine-readable benefit fragments
- glossary-aligned definitions
- entity-level metadata
- trust-layer signals
- validation-backed truth surfaces

Plan-ID is not the shift.

### **Structure is the shift.**

Plan-ID simply makes the shift impossible to ignore.

Plan-ID makes the shift visible, but the implications reach far beyond plan pages.

Every Medicare content type — benefit explanations, drug coverage descriptions, glossary terms, comparison narratives, county details, and star-rating summaries — now operates inside the same structured environment.

AI systems evaluate publishers not by keywords or page hierarchy, but by the consistency, clarity, and machine readability of the information they supply.

This is why **organizations that adopt structured publishing — not just carriers — will regain visibility across 2026–2028**, and those that do not will be represented by third parties, inferred data, or model-generated approximations.

The remaining question is not whether to adapt, but how.

Part III outlines the architecture Medicare publishers must implement to align with the AI retrieval ecosystem and rebuild durable Medicare visibility.

# PART III — THE NEW MEDICARE VISIBILITY ARCHITECTURE (THE FIX)

*The structural blueprint Medicare publishers must implement to regain plan visibility from AI Overviews, aggregators, and machine retrieval systems.*

The Medicare visibility transition is not reversible through traditional marketing levers. CMOs cannot fix this with more content, more paid search, more plan pages, or more brand spend.

The only sustainable path forward is structural.

Medicare organizations — including carriers, FMOs, agencies, and independent publishers — must adopt a visibility architecture that AI systems can ingest, trust, and reuse.

This requires alignment across:

- Plan-ID canonical identity
- machine-readable benefit and cost-sharing fragments
- glossary and terminology coherence
- location and market segmentation
- validation and governance workflows

Throughout 2024–2025, many organizations attempted to counter declining visibility with incremental SEO or additional content volume, not realizing that the underlying issue was architectural rather than tactical. **The Medicare Digital Architecture (MDA) reframes the challenge by addressing the structural elements AI systems depend on for accurate retrieval — regardless of who publishes them.**

PART III introduces the Medicare Digital Architecture (MDA) — the first industry-wide framework for ensuring consistent, accurate, AI-ready Medicare information.

This section answers: **What does the fix look like, and how must the Medicare ecosystem rebuild authority inside AI systems?**

## Chapter 11.

### The Plan-ID Entity Control Framework

#### ***How Medicare publishers regain visibility across AI systems.***

Plan-ID is the anchor of AI retrieval — but it is not the solution by itself.

It is simply the first structural point where Medicare publishers can regain control of how plan entities are interpreted inside AI systems.

Plan-ID is the “primary key” for:

- benefits
- counties
- networks
- plan comparisons
- formulary mapping
- star ratings
- out-of-pocket structures

AI systems already treat it this way.

Legacy aggregators dominate Plan-ID visibility today not because they publish structured data, but because they publish large volumes of predictable, plan-specific pages that models use as fallback sources when no machine-readable Plan-ID identity exists from the publisher.

Many organizations initially viewed Plan-ID displacement as a search-ranking issue rather than a structural one, not realizing that Plan-ID retrieval had moved inside AI systems — where structured publishing, not page hierarchy, determines which source becomes authoritative.

The Plan-ID Entity Control Framework establishes the first layer of structured authority by ensuring that Medicare plan identity is represented consistently, unambiguously, and in a machine-readable form that AI systems can trust.

## 1. Plan-ID Canonical Pages

Every plan must have a machine-readable canonical destination tied to:

- H-code
- Segment
- State
- county availability
- plan variation

These pages must be:

- fragmentable
- structured
- schema-anchored
- glossary-synchronized
- provenance-marked
- continuously updated

## 2. Plan-ID Redirect Logic

Publishers must unify authority signals by redirecting branded URLs to canonical Plan-ID endpoints.

This consolidates plan identity and prevents fragmentation across URLs.

## 3. Plan-ID Knowledge Panels

Structured data must enable Google, Bing, and AI systems to attach:

- ratings
- benefits
- summaries
- counties
- formulary attributes

...to the correct plan identity, rather than inferring or guessing.

## 4. Plan-ID Retrieval Surface

Each plan identity must be:

- indexable
- extractable
- queryable
- unambiguous

This is the minimum structural requirement for re-entering AI retrieval loops.

## Outcome of the Framework

Without this framework, AI retrieves fallback surfaces — aggregator pages, outdated CMS text, or inferred plan constructs — instead of authoritative, structured representations.

With this framework, **any Medicare publisher** that provides consistent, machine-readable Plan-ID identity becomes the preferred retrieval source, establishing the structural identity layer required for AI compatibility.

## Chapter 12.

### Structured Trust Surfaces for Medicare

***Why Medicare organizations must publish machine-readable benefits, fragments, and contextual explanations.***

AI systems do not trust pages.

They trust structured information — discrete, machine-readable fragments of truth that models can retrieve, recombine, and reuse consistently.

In Medicare, these trust surfaces fall into two categories:

1. plan-level truth (what CMS validates but does not publish), and
2. contextual truth (what CMS never intended to explain).

**Medicare publishers must produce both.**

Many organizations initially attempted to counter visibility decline by expanding narrative benefit pages or producing more long-form content, not realizing that AI systems require structured, fragment-level representations of plan and coverage truth rather than prose-based explanations. These misunderstandings left models to infer missing details, widening structural gaps in accuracy and visibility.

### Plan-Level Trust Surfaces: The Truth AI Cannot Infer

AI systems require plan information that is:

- explicit
- structured
- fragmentable
- year-specific
- county-bound
- glossary-aligned
- provenance-backed

Because CMS plan data is delivered through late-rendered React workflows and not exposed as machine-readable objects, AI systems cannot retrieve:

- detailed benefits
- cost-sharing rules
- network distinctions
- formulary tier structures
- county availability
- out-of-pocket patterns

AI reconstructs these from whoever publishes them — carriers, FMOs, agencies, aggregators, or independent Medicare publishers.

To become authoritative, Medicare publishers must provide:

- **structured benefit fragments**
- **structured cost-sharing fragments**
- **structured formulary and tier fragments**
- **structured metadata (plan type, segment, SNP classification, FIPS-bound counties)**
- **structured, year-bounded truth**
- **provenance metadata**

These surfaces anchor AI's understanding of the plan.

Without them, models infer — and hallucination fills the gaps.

## Sidebar: Why AI Requires Fragmentation

*How vector retrieval changes Medicare publishing*

AI systems do not retrieve content the way search engines do.

They do not index pages, scan HTML markup, or evaluate keyword density.

Instead, modern models use **vector retrieval** — a technique that represents information as mathematical embeddings that encode meaning, not text layout.

This has four implications:

### **1. Pages are too large to be represented as a single meaning.**

A plan page contains dozens of unrelated concepts:

copays, benefits, counties, networks, tiers, eligibility, disclaimers.

AI collapses these into a single averaged embedding, losing detail and often mixing concepts incorrectly.

## **2. AI retrieves the *closest* fragment of meaning — not the closest page.**

When someone asks:

“PCP copay for H5216-153?”

The model looks for a fragment representing **H5216-153 → PCP copay**, not a long-form page with 40 other attributes embedded in the same vector.

## **3. Fragment-level structure prevents hallucination.**

If each benefit, cost share, county binding, and tier rule exists as a discrete machine-readable object:

- AI retrieves only the relevant fragment
- the model does not infer missing data
- year-mixing and benefit-mixing collapse
- comparisons become stable and accurate
- plan identity becomes durable

## **4. Fragmentation is how AI “sees” Medicare.**

Humans read pages.

AI systems retrieve vectors.

Fragmentation translates Medicare information into the units AI needs — **entity-level, benefit-level, year-specific, county-bound objects.**

This is why structured publishing works:

it aligns Medicare content with the retrieval system AI actually uses.

Without fragmentation, models infer missing pieces from unrelated sources, creating the hallucinations, misclassifications, and compliance risks documented in this report.

## **Contextual Trust Surfaces: The Truth CMS Never Provides**

CMS validates benefits.

It does not explain Medicare.

AI systems must answer thousands of contextual questions CMS never addresses:

- “Does Medicare cover wheelchair ramps?”
- “Does Medicare cover glucose monitors?”

- “Does Medicare pay for transportation to appointments?”
- “Does Medicare help with dental implants?”

**Medicare publishers who provide clear, contextual, glossary-aligned explanations create trust surfaces AI relies on to interpret coverage rules.**

Humana is one such example — and the effect is immediate:

- AI sees Humana as knowledgeable
- Their explanations inform AI’s interpretation of benefit categories
- Their content clarifies distinctions between Original Medicare and supplemental benefits

**The Medicare.org rebuild further validated this principle:**

Its 64 structured, county-level contextual insights — premiums, plan type distributions, dental prevalence, \$0-premium penetration, star concentration — provided AI with the local context CMS never provides but that users need.

When Medicare organizations do not publish contextual surfaces:

- AI fills gaps with outdated CMS text
- or relies on aggregator interpretations
- or invents its own rules

Context is not optional.

It is a structural requirement for AI visibility.

## The Dual Publishing Mandate: Humans and AI Must Be Served Together

In the pre-AI web, publishers wrote for:

- consumers
- agents
- SEO

In the AI era, publishers must write for:

- humans **and**
- AI systems

This requires:

- readable human explanations

- structured fragments for AI
- glossary coherence for both
- county-level context for relevance
- metadata for retrieval
- narratives that do not conflict with structured truth

AI cannot aggregate what Medicare publishers do not supply.

AI cannot explain what publishers do not define.

When publishers produce structured surfaces for both audiences, AI systems stabilize — and the publisher becomes the default answer.

## Structured Trust Surfaces Lead Directly to AI Authority

Structured plan-level surfaces combined with structured contextual surfaces allow AI systems to:

- resolve plan identity
- understand benefit categories
- compare plan features
- differentiate plans correctly
- minimize hallucinations
- connect user intent to plan logic
- generate compliant explanations
- recommend plans more reliably

This is visibility in the AI era.

Not rankings.

Not clicks.

Not keywords.

### **Structured trust surfaces.**

As AI becomes the primary gateway to Medicare plan understanding, the ecosystem will require a unified, structured “Medicare Trust Surface” — a consistent reference layer that resolves plan identity, benefit definitions, county context, and coverage logic without ambiguity. Because CMS does not expose machine-readable interpretive structures, and because Medicare publishers produce heterogeneous formats, this layer will need to be maintained by a neutral, nationally connected Medicare publisher capable of serving all carriers without brand bias.

## Chapter 13.

### Glossary Alignment (The Medicare Terminology Crisis)

***AI is hallucinating definitions. Medicare organizations must standardize Medicare terms or risk misrepresentation.***

AI systems hallucinate most often when terminology is:

- ambiguous
- inconsistent
- outdated
- contradictory across sources
- or missing entirely

Medicare terminology suffers from all of these issues:

- CMS definitions are regulatory, not explanatory — and not published in a structured, machine-readable form
- Medicare publishers use terms inconsistently
- similar benefits are named differently across plans
- internal plan constructs rarely match public descriptions
- no standardized, machine-readable glossary exists for AI systems

Because early visibility declines were attributed to algorithm volatility or SEO variation, glossary work was often deprioritized or treated as a content refinement task. In reality, terminology alignment is a structural requirement: AI systems rely on consistent definitions to interpret benefit categories, supplemental services, and coverage rules without invention or drift.

This creates:

- conflicting Medicare definitions

- hallucinated coverage rules
- incorrect benefit interpretations
- misclassification across carriers
- inaccurate plan comparisons
- compliance exposure
- enrollment risk

The industry currently treats “glossary work” as marketing. In the AI era, it is **infrastructure**.

Glossary alignment is also essential for fragment-level retrieval, because AI systems anchor benefit fragments to vocabulary patterns. Inconsistent terms produce inconsistent embeddings, increasing the likelihood of hallucination.

To stabilize AI explanations and prevent the model from inventing definitions, carriers need:

- Clear, consistent terminology
- Definitions aligned with CMS but adapted for AI interpretation
- Uniform usage across plan pages, benefits, and county content
- Structures that ensure models interpret benefit terms correctly

If Medicare organizations do not define their terms clearly and consistently, AI systems will supply definitions themselves — and those definitions may be inaccurate.

In an AI-first ecosystem, glossary alignment is a shared responsibility. The Medicare ecosystem will ultimately require a unified, structured vocabulary — a trust surface for terminology — to prevent drift, stabilize interpretation, and ensure consistent benefit logic across carriers, markets, and publishers.

## Chapter 14.

### Machine-Readable Plan and Location Content

*The migration from narrative Medicare pages to structured data surfaces.*

AI systems cannot reliably interpret Medicare information from narrative plan pages. They require structured, consistent, machine-readable representations of:

1. **What the plan is**  
(its identity, benefits, cost-sharing, metadata)
2. **Where the plan exists**  
(county availability, FIPS alignment, geography)
3. **Who the plan serves**  
(market segment: MA/MAPD, D-SNP, C-SNP, I-SNP, and other population-specific categories)

Plans are not abstract objects.

They exist in specific locations and serve specific populations with different eligibility, needs, and benefit priorities.

AI retrieval depends on all three dimensions.

Because early visibility loss appeared similar to SEO decline, many organizations continued producing narrative plan content without realizing that AI systems no longer interpret those pages directly. This prolonged reliance on prose-based plan descriptions widened the structural gap and left models to infer missing details across plan, location, and market dimensions.

### Plan-Level Structure: What the Plan Is

Every Medicare Advantage plan must be published by **the plan owner (the carrier)** as a coherent, machine-readable identity.

This includes structured representations of:

- plan type and segment
- star rating
- SNP classification
- cost-sharing rules
- benefits and supplemental benefits
- drug tiers and network qualifiers
- metadata and compliance fields
- yearly distinctions and effective dates

Narrative plan pages — mixed with marketing language, disclaimers, and cross-year phrasing — do not supply the clarity AI requires.

Structured plan content allows AI systems to:

- retrieve the correct plan
- differentiate similar plans
- avoid merging H-codes
- reduce benefit hallucinations
- compare plan features accurately
- bind explanations to the correct plan year

Without structured plan content, AI fills gaps with inference — and inference produces error.

## Location-Level Structure: Where the Plan Exists

CMS does not publish county-level Medicare insights for AI systems.

Yet AI must answer inherently geographic questions:

- “Which plans are available in my county?”
- “How many PPOs are in Hillsborough County?”
- “What is the average MOOP in Bernalillo County?”

County-level information is critical because:

- availability varies
- networks vary
- plan density varies
- competition varies
- premiums vary

- benefits vary
- member experience varies

To support accurate AI retrieval, **Medicare publishers** — including carriers and independent entities — must publish structured location content that includes:

- service area binding
- FIPS-driven county identity
- normalized state/county naming
- year-specific availability
- consistent metadata
- clear county-to-plan relationships

This is the foundation for geographic reasoning inside AI systems.

## Market-Level Structure: Who the Plan Serves

CMS does not publish an explicit, machine-readable market segmentation model. Plans serving fundamentally different populations — MA/MAPD, D-SNP, C-SNP, I-SNP — are often presented side-by-side in digital plan catalogs, even though they serve different eligibility groups and clinical or socio-economic needs.

AI systems already infer segmentation through language and context:

- Medicaid eligibility
- chronic conditions
- institutional status
- disability
- dual eligibility
- condition-specific qualifiers
- care coordination requirements

But without structured market surfaces, these inferences remain inconsistent.

To support accurate segmentation, **Medicare publishers** must publish:

- SNP vs MA/MAPD eligibility distinctions
- D-SNP Medicaid-bindings
- C-SNP condition qualifiers
- I-SNP institutional context
- plan-specific population intent

- contextual explanations aligned with structured data

Without this segmentation layer, AI may misclassify plans, compare incompatible designs, or recommend plans to ineligible beneficiaries.

## Contextual County-Level Truth: The Data CMS Never Provides

AI must answer user questions CMS never attempts to explain.

This is the role of derived county-level context — the type of contextual intelligence demonstrated in the Medicare.org rebuild:

- total plan counts
- HMO–PPO distribution
- \$0-premium prevalence
- average premiums
- average MOOP
- star rating concentration
- dental/vision/hearing prevalence
- year-over-year churn
- benefit availability patterns

These insights act as **contextual trust surfaces**, allowing AI to:

- reason locally
- summarize markets accurately
- frame comparisons
- reduce geography-driven hallucinations

This is information CMS never intended to publish — and exactly what AI systems require.

## Machine-Ready Explanations for Humans and AI

Medicare information must serve two audiences at once:

1. humans reading the content
2. AI systems retrieving the content

This requires:

- fact-first narratives
- glossary alignment
- year specificity
- county correctness
- segment accuracy
- structured fragments
- CMS-compliant logic

When narrative explanations conflict with structured truth, AI privileges the structure — or infers incorrectly.

## Provenance and Versioning: The Trust Backbone

AI systems rely on consistency and transparency.

Each plan, location, and market surface must include:

- Source
- effective year
- last updated
- version
- timestamp

Provenance stabilizes retrieval and signals that the publisher maintains an authoritative record over time.

## Why This Matters

Without machine-readable plan, location, and market content:

- AI misassigns counties
- plan features blur
- SNP and MA plans conflate
- supplemental benefits disappear
- cost-sharing gets inferred incorrectly
- eligibility rules become unstable
- models mix plan-years
- aggregator fallbacks dominate
- hallucinations multiply

With structured content:

- plan identity stabilizes
- AI explanations become accurate
- county reasoning improves
- market segmentation clarifies
- supplemental benefits surface correctly
- comparisons align with reality
- structured truth becomes the retrieval foundation
- **the publishing entity becomes the primary structured source AI trusts**

This is not a design decision.

**It is the structural requirement for Medicare visibility in 2026–2028.**

## Chapter 15.

### The AI Validation Layer

#### ***Carrier compliance in the generative era.***

As Medicare organizations adopt structured plan, location, and market content, a new challenge emerges: ensuring that AI systems interpret, reuse, and represent that content accurately over time. Generative AI introduces a dynamic layer of risk.

Models:

- reinterpret content,
- infer missing details,
- synthesize explanations,
- generalize across carriers,
- and drift as they update.

Because early AI misstatements were often seen as generic “hallucinations,” rather than the predictable result of models reinterpreting incomplete or inconsistent plan data, many organizations underestimated the need for a dedicated validation function even after adopting structured publishing.

Structured publishing reduces hallucinations but cannot eliminate them.

Without a formal validation layer, Medicare publishers face:

- cross-year inconsistencies,
- plan misrepresentation,
- incorrect comparisons,
- benefit distortion,
- network misassignment,
- glossary drift,

- and widening compliance exposure.

AI validation becomes a **core Medicare function** — not optional, not tactical, and not a marketing activity.

It becomes an **enterprise governance requirement**.

## The Three Capabilities of the AI Validation Layer

### *1. Accuracy Monitoring*

Medicare organizations must continuously monitor how AI systems describe their plans and contextual explanations across:

- GPT
- Gemini
- Perplexity
- AI-powered browsers
- downstream agentic tools

This identifies where model interpretations diverge from MDA-aligned structured truth.

### *2. Consistency Enforcement*

Validation must ensure:

- plan-year boundaries remain intact
- glossary terms remain consistent
- county information remains correct
- benefit definitions remain aligned
- market segmentation remains explicit
- Plan-ID identity remains stable

Without this discipline, LLMs drift — and internal content drifts with them.

### *3. Pre-Publication & Post-Publication Verification*

All AI-generated or AI-assisted explanations — internal and external — must be validated:

- before use in agent workflows
- before appearing on carrier websites
- before being shared with members
- before influencing decision support tools

This ensures that AI does not introduce unapproved language, outdated values, or incorrect benefit logic.

The AI Validation Layer is not SEO.

It is not editorial QA.

It is **enterprise governance for AI-driven Medicare information.**

## Why Medicare Organizations Cannot Avoid Validation

As generative systems increasingly become the first point of plan discovery, publishers without an AI validation capability will face:

- recurring misstatements
- consumer confusion
- agent-side misinformation
- unintended plan steering
- widening compliance risk
- CMS audit exposure
- enrollment distortions
- legal liability

Organizations that implement validation will ensure their structured data is:

- interpreted correctly
- retrieved consistently
- represented accurately
- compliant across surfaces
- reinforced as the authoritative source

The Validation Layer is the **operational safeguard** for Medicare visibility and accuracy in 2026–2028.

It keeps structured truth aligned with how AI systems actually describe Medicare content.

## Sidebar: AI Browser Governance & Shadow AI Risk

*Why Medicare organizations must control AI browser usage before it becomes a compliance liability.*

As AI browsers mature, many teams experiment with consumer-grade AI extensions for summarization, comparisons, or compliance checks. These tools provide convenience — but represent **major enterprise risks** when used without governance.

### Shadow AI Risk

Unvetted AI browser extensions send page content to external APIs, creating exposure to:

- unauthorized data transmission
- loss of audit trails
- unapproved model usage
- prompt injection attacks
- inadvertent disclosure of proprietary or PHI-related information

These behaviors bypass IT controls and are invisible to compliance teams.

### MDA-Governed AI Browser Usage

To use AI browsers safely, carriers must ensure extensions:

1. operate in a controlled, enterprise environment;
2. ground truth exclusively in MDA-aligned structured plan data;
3. log all interactions for compliance review;
4. validate year, county, and eligibility alignment;
5. enforce CMS-compliant terminology and benefit definitions.

With MDA as the authoritative truth source, AI browsers shift from compliance liabilities to **real-time quality gates** that catch misstatements before they reach consumers or CMS.

## PART IV — THE 2026–2028 OUTLOOK (THE FORECAST)

*The future of Medicare visibility, growth, and digital competition in an AI-driven environment.*

As AI systems continue to intermediate plan discovery, Medicare visibility will consolidate around structured data sources, validated plan identity, and consistent glossary alignment.

**Organizations that modernize early — regardless of role — will gain durable advantages that compound across enrollment seasons. Those that delay will be structurally displaced not by competitors directly, but by AI systems that rely on the structured sources available to them.**

By 2028, Medicare visibility will not be determined by SEO, branding, or competitive spend, but by whether an organization has become a structured source of truth. Those that do not adopt structured approaches may experience material reduction in AI-mediated discovery visibility.

Because many teams initially interpreted visibility decline as an SEO or traffic fluctuation rather than a structural realignment, the competitive gap widened during 2024–2025 — and those early divergences will accelerate through 2028 as AI retrieval stabilizes around consistent sources of structured truth.

PART IV outlines:

- who gains visibility and who loses it
- how compliance risk escalates
- how AI will shape enrollment dynamics
- how Plan-ID retrieval becomes the dominant discovery vector
- how aggregators decline
- how FMOs evolve
- how carrier brands shrink or grow based on structured publishing
- how AI agents become the primary comparison layer

This section answers: **What will the Medicare digital landscape look like in 2028 — and who is positioned to win?**

## Chapter 16.

### Structural Advantages and Emerging Risks (2026–2028)

***Which organizations are positioned to gain visibility—and which face increasing structural challenges.***

This forecast is based on:

- Plan-ID surface analysis
- structured publishing readiness
- benefit fragmentation maturity
- glossary and terminology consistency
- metadata coherence
- AI Overview inclusion
- retrieval stability across GPT, Gemini, and Perplexity
- county-level accuracy trends
- content architecture depth

Because much of the industry initially treated visibility loss as an SEO or competitive ranking issue, organizations that delayed structural adaptation now face a widening gap; the advantages and disadvantages outlined here reflect that early divergence and will accelerate as AI retrieval stabilizes around structured authority sources.

### Organizations Structurally Positioned for Advantage

#### **1. Organizations with Structured Plan Identity**

Organizations that adopt machine-readable Plan-ID surfaces early will see improved retrieval stability, fewer hallucinations, and higher representation in AI-generated summaries. Structured identity becomes a durable advantage by 2027.

## **2. Publishers with Fragment-Level Content**

Entities that provide multi-fragment plan representations—benefits, cost sharing, drug tiers, county bindings—will become preferred sources for AI systems interpreting Medicare data.

## **3. Organizations that Build Trust Surfaces Early**

Structured contextual content (“Does Medicare cover...?”), year-specific explanations, and glossary-aligned narratives help AI systems distinguish plan benefits and interpret complex benefit categories. Early adopters will become recurring citation surfaces across models.

## **4. FMOs Adopting AI-Governed Comparison Tools**

FMOs that modernize agent workflows with structured plan comparisons, county-specific insights, and validation references will outperform organizations relying on static PDFs or CRM notes.

## Organizations Structurally Disadvantaged

### **1. Carriers with Brand-Heavy Naming Conventions**

Plan families with similar or overlapping names create entity confusion. AI systems rely on structured identity, not branding, making traditional naming strategies less effective.

### **2. Aggregators Lacking Structured Architectures**

Thin comparison pages and SEO-oriented templates will lose visibility as AI systems increasingly generate plan comparisons internally. Models will prefer structured fragments over narrative summaries.

### **3. Agencies Relying on Narrative, SEO-Era Content**

AI systems bypass unstructured content entirely. Without structured surfaces, agencies will struggle to influence plan interpretation or downstream explanations.

### **4. Carriers without glossary alignment**

Inconsistent terminology leads to increased model inference, generating benefit hallucinations and incorrect explanations over time. This becomes a growing compliance risk through 2027–2028.

## Functions Likely to be Replaced or Absorbed by AI Systems

- generalized “Top 10” Medicare plan lists
- coverage and FAQ explainer pages

- basic plan comparison surfaces
- unstructured benefit summaries

AI systems will increasingly synthesize these functions independently, reducing traffic to traditional digital surfaces.

## Chapter 17.

### The Medicare Plan-ID Visibility Index (MVI)

*An industry-level measurement of Medicare visibility in the AI era.*

The Medicare Plan-ID Visibility Index (MVI) is an annual benchmark that evaluates **plan-level and publisher-level visibility** across both search and AI-driven retrieval surfaces. It measures visibility for specific plans and their associated publishers across:

- Google
- AI Overviews (AIO)
- GPT retrieval
- Gemini retrieval
- Perplexity retrieval
- aggregator presence
- independent publishers

As traditional SEO metrics became less predictive of Medicare visibility during 2024–2025, organizations needed a way to measure performance based on structural signals rather than keyword rankings or traffic patterns — MVI fills that gap by evaluating the elements AI systems actually use to interpret and surface plans.

The Index evaluates Medicare publishers using five weighted dimensions:

1. **Plan-ID Retrieval Strength (40%)**

The reliability with which AI systems surface a carrier’s plans when interpreting plan-level or benefit-level queries.

2. **Structured Authority Score (25%)**

The presence of machine-readable benefit structures, glossary-aligned terminology, and plan metadata.

3. **AIO Presence Score (15%)**

The frequency and clarity of a carrier's plan representation in Google AI Overviews.

4. **Entity Clarity Score (10%)**

How easily AI systems distinguish one plan from another across a carrier's portfolio.

5. **Fragment Stability Score (10%)**

The consistency of plan explanations across multiple AI models, years, and county contexts.

The MVI provides carriers, FMOs, agencies, and independent Medicare publishers with an objective visibility benchmark as Medicare transitions from SEO-based discovery to AI-driven retrieval.

## Chapter 18.

### The 2026–2028 Medicare AI Risk Map

*A strategic forecast of accuracy, compliance, retrieval drift, and digital exposure.*

The following risk horizon reflects expected AI-driven pressures on Medicare information systems. Because much of the industry initially addressed visibility decline through traditional marketing or SEO adjustments rather than structural alignment, these risks intensify along a predictable trajectory as AI systems consolidate around consistent sources of machine-readable truth.

A key reality underlies all three years of this forecast:

**AI systems are not choosing winners — they are choosing structures. Organizations without structured Medicare truth may experience significantly reduced representation in AI-mediated retrieval environments.**

#### 2026: The Model Interpretation Year

- AI systems increase reliance on structured surfaces
- unstructured pages experience visibility suppression
- cross-year hallucinations increase without glossary alignment
- publishers experience accuracy variance across GPT and Gemini updates
- CMS may issue new AI-related guidance

**Risk level: High**

#### 2027: The Structural Displacement Year

- Plan-ID retrieval becomes the dominant plan discovery vector

- aggregator visibility decreases significantly
- AI assistants begin driving more comparison behavior than websites
- carrier visibility increasingly differentiates based on structural readiness
- FMOs modernize agent workflows using AI-interpreted plan data

**Risk level: Very High**

## 2028: The Structural Consolidation Year

- visibility advantages built through structured publishing become self-reinforcing
- carriers lacking structured surfaces experience reduced inclusion within AI-generated comparison flows
- AI-driven plan interpretation shapes enrollment patterns
- agencies without AI-aligned materials lose competitive positioning
- FMOs consolidate around data-driven, structured models

**Risk level: Elevated and systemic**

**Opportunity level: Significant for early adopters**

## Chapter 19.

### The Medicare Digital Displacement Curve

*How AI reshapes the Medicare digital ecosystem through 2026–2028.*

The Medicare Digital Displacement Curve illustrates four concurrent shifts:

#### **1. Aggregators → AI Agents**

As AI systems answer benefit questions, perform plan comparisons, and surface Plan-IDs, aggregator visibility naturally declines.

#### **2. Narrative Websites → Structured Surfaces**

Narrative plan pages and SEO-era content lose visibility.

Structured benefit fragments, metadata, and county-level truth surfaces become dominant retrieval signals.

#### **3. Legacy Agencies → AI-Augmented Agencies**

Agencies dependent on narrative explanations and static comparison tools see declining accuracy and retention.

AI-assisted, data-driven workflows become the norm.

#### **4. FMOs → Tech-Integrated FMOs**

FMOs evolve into structured data hubs, agent-support platforms, and content governance nodes.

Those who adapt gain influence within the Medicare digital supply chain.

Because early visibility shifts resembled typical SEO or channel volatility, many organizations underestimated how quickly AI systems were replacing entire layers of digital interaction. These structural transitions compound over time, driving the displacement patterns shown in the Digital Displacement Curve.

## THE NEW MEDICARE POWER MAP (2028)

By 2028, Medicare visibility will be shaped by:

- publishers providing structured, machine-readable plan and contextual content
- carriers that expose consistent Plan-ID identity
- FMOs with AI-aligned agent workflows
- agencies equipped with structured comparison surfaces
- AI systems with stable retrieval pathways

Not by:

- SEO
- legacy websites
- unstructured content
- brand recognition alone

The structural shifts outlined here are already underway and will accelerate through 2028.

Organizations that adopt structured publishing early will hold a durable visibility foundation.

Those that delay will face increasing difficulty recovering influence within AI-based discovery systems.

# PART V — The Medicare Digital Architecture (MDA)

*A foundational framework for structured Medicare information in the AI era.*

The structural changes outlined in this report point to a clear conclusion:

**Medicare requires a new digital framework—one that supports machine-readable plan identity, accurate interpretation, and consistent retrieval across AI systems.**

The Medicare Digital Architecture (MDA) is the conceptual model that defines this foundation.

It is not a software platform, a marketing tactic, or a publishing strategy.

**It is the underlying information architecture AI systems require in order to represent Medicare plans responsibly and accurately.**

In an AI-first Medicare environment, the authority of any individual plan page or benefit explanation depends on its alignment with a broader ecosystem of structured truth. Over time, this ecosystem will consolidate around one or more neutral publishers that operate across carriers, markets, and channels — effectively forming a **Medicare Trust Surface** that AI systems rely on as the canonical reference layer.

MDA aligns five critical components of Medicare information:

## **1. Structured Plan Identity**

Plan-ID anchored representations of each plan, including:

- plan metadata
- benefits
- cost sharing
- network structure
- drug tiers

- county-level availability
- year-specific versions
- provenance signals

This ensures models can consistently retrieve and differentiate plans.

## **2. Structured Location Context**

County-level surfaces that include:

- availability
- network/geography context
- benefit variation
- competitive landscape
- derived county insights

AI must interpret Medicare **in location**, not in abstraction.

## **3. Glossary and Terminology Alignment**

Consistent Medicare definitions—aligned with CMS but modernized for AI interpretation—so models do not infer or hallucinate benefit rules.

Terminology coherence stabilizes plan explanations and supports multi-model accuracy.

## **4. Structured Contextual Content**

Human-readable and machine-readable explanations that address real beneficiary questions:

- “Does Medicare cover...?”
- condition-specific benefits
- ancillary services
- home-based supports
- chronic condition needs

This contextual truth is the information layer CMS does not provide but AI depends on.

## 5. Validation and Governance

A continuous process that ensures:

- structured content remains accurate
- plan-year distinctions stay intact
- glossary terms remain consistent
- county bindings do not drift
- AI-generated explanations reflect actual plan design

MDA positions validation as an enterprise responsibility, not a marketing task.

### Why MDA Matters

MDA provides the industry with a unified, structural framework for publishing Medicare information that AI systems can retrieve, interpret, and reuse accurately.

By organizing Medicare plan information into a consistent architecture — plan identity, location context, terminology coherence, contextual explanations, and validation — MDA:

- reduces hallucinations
- improves accuracy
- restores carrier authority
- stabilizes plan comparisons
- supports compliance
- enhances beneficiary trust
- and creates a durable visibility foundation

AI does not require more content.

AI requires **structured truth**.

MDA defines the framework for delivering it.

## Who MDA Serves

MDA supports:

- **Carriers** seeking visibility, accuracy, and compliance
- **FMOs** aligning agent tools with AI-driven interpretation
- **Agencies** modernizing digital Medicare publishing
- **AI systems** requiring structured Medicare knowledge
- **Beneficiaries** relying on accurate, consistent plan explanations

## The Role of a Medicare Digital Standards Body

As the Medicare ecosystem transitions into an AI-first discovery environment, long-term stability may require the emergence of a neutral, standards-aligned trust layer capable of producing consistent, machine-readable representations of plan identity, benefit definitions, terminology, and geographic context across carriers and markets.

This report does not designate any specific entity to fulfill that role. The Trust Publishing Institute (TPI) operates as a time-bound research body studying these structural conditions and documenting architectural characteristics correlated with retrieval stability.

# Epilogue

## *The Work Ahead*

Medicare’s digital transformation is no longer theoretical.

It is underway — accelerated by AI systems that interpret and communicate plan information as part of a new, machine-driven discovery layer.

This shift has revealed structural gaps the industry has never been required to resolve: gaps in plan identity, benefit definitions, location context, terminology, and year-specific truth.

Carriers, FMOs, agencies, independent publishers, and digital teams now operate in an environment where accuracy, visibility, and trust depend on the quality and structure of the information they publish. The transition is not optional. It is foundational to the future of beneficiary understanding and plan selection.

This report outlines a framework for navigating that transition:

- a move from narrative pages to structured surfaces;
- from keyword-driven content to entity-driven architecture;
- from static plan explanations to machine-readable truth;
- from isolated plan pages to a coherent Medicare entity graph.

The ecosystem that emerges through 2026–2028 will reward clarity, structure, consistency, provenance, and validation.

It will favor publishers who act as stewards of accurate Medicare information.

And it will increasingly rely on organizations that maintain high-integrity digital representations of their plans, benefits, and counties.

The Medicare ecosystem will not stabilize until a consistent, machine-readable trust layer exists — one that unifies plan identity, coverage definitions, county reality, glossary coherence, and structured benefit truth. As the industry transitions into AI-mediated discovery, this trust surface will almost certainly emerge from a neutral, nationally connected Medicare publisher or standards body with the infrastructure, relationships, and domain footprint required to maintain it at scale.

The path forward requires collaboration between marketing, technology, compliance, clinical operations, and enterprise strategy — and a shared recognition that Medicare digital governance is now an organizational responsibility. Early adopters will not simply restore visibility; they will help shape the standards that define how Medicare information is retrieved, interpreted, and trusted in the years ahead.

**A neutral Medicare standards body — whether an institute, consortium, or multi-stakeholder governance group — will be needed to guide this transition, maintain structured terminology, align plan identity conventions, and ensure responsible, accurate representation across AI-driven systems.**

The landscape has changed.

The work continues.

## About the Author

**David Bynon is a Medicare data-architecture specialist whose work focuses on structured Medicare publishing, plan-data integrity, AI-ready content systems, and the emerging dynamics of AI-driven visibility in regulated healthcare environments.**

During the rollout of Google AI Overviews, Bynon rebuilt the Medicare.org platform using machine-readable plan identity, structured benefit fragments, and county-level contextual surfaces — resulting in the first documented visibility recovery under AI conditions. This work provided real-world empirical validation for the structured publishing principles outlined in this report.

Bynon's background includes enterprise systems design, large-scale data architecture, high-availability infrastructure, and more than twenty years of structured publishing innovation across regulated industries. His research into AI retrieval behavior, Plan-ID entity alignment, benefit-level fragmentation, and structured trust surfaces forms the foundation of the Medicare Digital Architecture (MDA) framework.